arXiv:1611.04786v1 [cs.CR] 15 Nov 2016

AdversariaLib: An Open-source Library for the Security Evaluation of Machine Learning Algorithms Under Attack

Igino Corona
Battista Biggio
Davide Maiorca
Dept. of Electrical and Electronic Engineering
University of Cagliari
Piazza d'Armi, 09123, Cagliari, Italy
{IGINO.CORONA,BATTISTA.BIGGIO,DAVIDE.MAIORCA}@DIEE.UNICA.IT

Editor:

Abstract

We present AdversariaLib, an open-source python library for the security evaluation of machine learning (ML) against carefully-targeted attacks. It supports the implementation of several attacks proposed thus far in the literature of adversarial learning, allows for the evaluation of a wide range of ML algorithms, runs on multiple platforms, and has multi-processing enabled. The library has a modular architecture that makes it easy to use and to extend by implementing novel attacks and countermeasures. It relies on other widely-used open-source ML libraries, including *scikit-learn* and *FANN*. Classification algorithms are implemented and optimized in C/C++, allowing for a fast evaluation of the simulated attacks. The package is distributed under the GNU General Public License v3, and it is available for download at http://sourceforge.net/projects/adversarialib.

Keywords: adversarial learning, security evaluation, evasion attacks

1. Introduction

Machine learning algorithms have been increasingly adopted in security applications like spam, intrusion and malware detection, mainly due to their ability to generalize, *i.e.*, to potentially detect novel attacks or variants of known ones. However, in adversarial settings such as the aforementioned ones, the underlying assumption of data *stationarity*, *i.e.*, that both training and test data are drawn from the same (though possibly unknown) distribution, is often violated by intelligent and adaptive adversaries that purposely manipulate data to compromise learning. This has recently raised the issue of evaluating the security of learning algorithms to carefully-targeted attacks, besides designing suitable countermeasures, as pointed out in the growing research area of *adversarial machine learning* (Barreno et al., 2006; Huang et al., 2011; Biggio et al., 2014).

In this work we present AdversariaLib, an open-source python library that can be exploited to this end. To the best of our knowledge, this is the first open-source library that implements the process of security evaluation recently advocated by Biggio et al. (2014). The idea behind a *security evaluation* is to *proactively* anticipate the behavior of the adversary to identify potential vulnerabilities of machine learning algorithms, and to

design suitable countermeasures, *before* the corresponding attacks may actually occur. To this end, potential attacks against the given classifier are *simulated* according to a given adversary's model through an ad hoc manipulation of training and test data, and then their impact on the targeted classifier's performance is evaluated. The core of this procedure is clearly the way in which the training and the test data are modified, depending on the considered attack. For instance, in the *evasion* setting, the adversary can only manipulate malicious samples (*e.g.*, spam emails) at test time to evade detection, while training data remains unchanged. How each sample is manipulated depends on specific assumptions made on the adversary's model, *i.e.*, on her goal, knowledge of the attacked system, and capabilities of manipulating the data (Biggio et al., 2014, 2013). It is thus clear that the attack strongly depends on the targeted classifier. In practice, however, only rarely the adversary may know the targeted classifier completely: although she may realistically know the kind of algorithm used, she may not know the parameters learned after training (*e.g.*, the weights assigned to each feature by a linear classifier).

Another interesting feature of `AdversariaLib` is indeed the possibility of learning an estimate of the targeted classifier, which we refer to as *surrogate* classifier. This assumes that the adversary can collect a set of surrogate training data, ideally drawn from the same distribution as the targeted classifier's training set, retrieve the labels assigned to them by the targeted classifier, and learn a *surrogate* classifier on such data. In Biggio et al. (2013) we showed that attacking a surrogate classifier learned on a relatively small training set may lead to evade the targeted classifier with high probability. This is an example of how machine learning can be offensively exploited to defeat machine learning-based defenses. `AdversariaLib` natively implements this mechanism to simulate *realistic* attacks in which the adversary has limited knowledge of the attacked system.

We present the library's architecture and implementation in Sect. 2, as well as an attack example in Sect. 3. Conclusions and future extensions are discussed in Sect. 4.

2. `AdversariaLib`: architecture and implementation

`AdversariaLib`, as shown in Fig. 1, is structured into three main modules. The first, `advlib`, implements the attack algorithms and functions for the evaluation of classifiers under attack. The second, `prlib`, implements standard functions for handling datasets, defining distances between samples in feature space, learning classifiers, and evaluating their performance. To this end, `prlib` implements suitable wrappers that exploit other open-source libraries, such as *scikit-learn* (Pedregosa et al., 2011) and *FANN*.[1] The third, `util`, contains library functions for storing datasets, classifiers, and results in the filesystem, and for logging the progress of current experiments (on the standard output).

The library's main script, named `runexp`, implements a rather general experimental scheme that can be used to evaluate the security of different machine learning algorithms against several kinds of adversarial attacks. In particular, it executes the following steps: (*i*) randomly split data into training-test pairs (if such pairs are not given initially); (*ii*) for each training portion, learn the corresponding classifier(s); (*iii*) attack each learned classifier; and (*iv*) store the results. Any experiment is executed according to the parameters specified

1. *FANN* is an open-source C library for learning Fast Artificial Neural Networks (not natively supported by *scikit-learn*), available at `http://leenissen.dk/fann/wp/`.

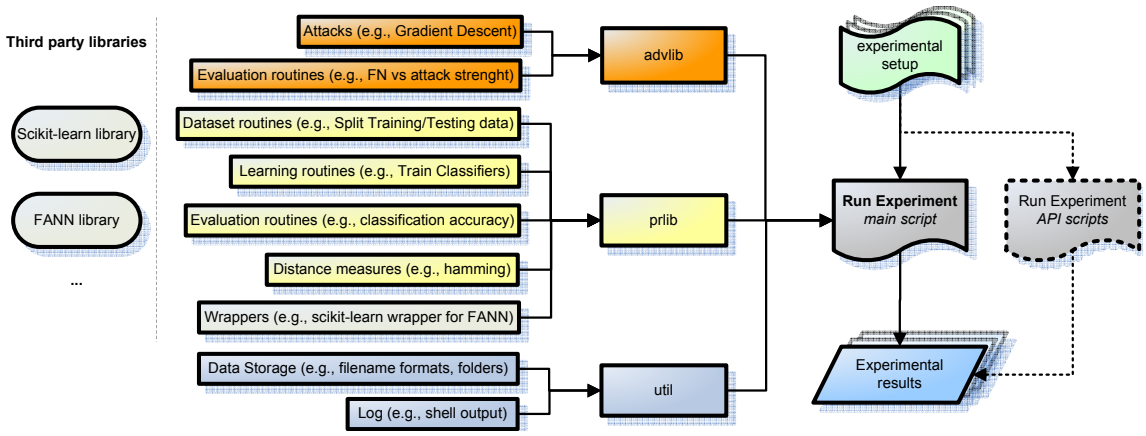

Figure 1: Architecture and main components of `AdversariaLib`.

in the *setup* file, such as, *e.g.*, number of training-test pairs, type and parameters of each targeted classifier, and type and parameters of the chosen attack(s). It is possible to run different experiments through the proper definition of different setup files. Experimental results are saved according to the format defined in the `util` module, and depending on the given attack. `AdversariaLib` also provides a set of *API* scripts to run simple experiments, such as training a classifier, attacking it, and evaluating its classification accuracy.

Implementation. `AdversariaLib`, as well as its dependencies (third-party libraries) is fully open-source, multi-platform, and can be used flawlessly in all major operating systems, since it is written mainly in Python. An important aspect is that Python allows researchers to easily and quickly implement new functionalities, and has a huge developers community. Moreover, an in-depth code optimization can be easily achieved exploiting C/C++ language bindings for Python, as performed by the external libraries *scikit-learn* and *FANN*. We carefully documented the whole library, with an extensive set of usage examples, as described in the official website `http://comsec.diee.unica.it/adversarialib`. We also provided a detailed documentation on how to implement and add new attacks to the library, to make it easy to extend. Each attack implementation must provide a simple, general interface, and the related files must be put in a separate folder within the `advlib` module (see Fig. 1).

FANN and Matlab wrappers. We implemented a scikit-learn wrapper for FANN classifiers (see `prlib/classifier/mlp.py`). Following the same approach, novel classifiers, *e.g.*, learning algorithms robust to adversarial attacks, can be easily added to the library. We also provide an open-source wrapper that allows us to configure and run experiments using `AdversariaLib` directly from the Matlab environment. In addition, the corresponding results can be exported as `PDF` files. This wrapper relies on the APIs of `AdversariaLib`, and can thus easily execute the evasion attack described by Biggio et al. (2013). An example is presented in the next section. This wrapper is freely downloadable from the official `AdversariaLib` repository.

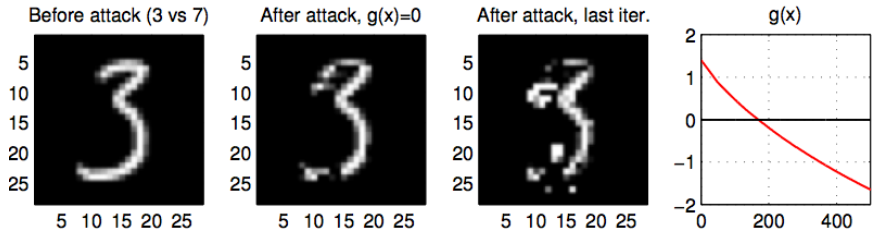

Figure 2: Gradient-based evasion of a linear SVM trained on handwritten digits.

3. Attack example on handwritten digit recognition

We report here a simple example to visually demonstrate the effectiveness of a gradient-based evasion attack, adapted from our recent work in Biggio et al. (2013). This example can be run using the Matlab wrapper of `AdversariaLib` and, in particular, the script `main_mnist.m`. A linear SVM is first trained to discriminate between two classes of handwritten digits from the MNIST dataset (LeCun et al., 1995), *i.e.*, `3` (positive class) and `7` (negative class), and then, the former is manipulated to evade detection. The attack iteratively modifies the initial `3` to minimize the SVM's discriminant function $g(\mathbf{x})$ through gradient descent. Results are shown in Fig. 2. The leftmost `3` represents the initial, non-manipulated attack sample. The second and third depicted digits represent *evasion points*, *i.e.*, manipulated `3`s that are misclassified as `7`, respectively as soon as the SVM's discriminant function $g(\mathbf{x})$ becomes non-positive (*i.e.*, the first found evasion point during the descent), and after more iterations. This example not only highlights that few manipulations are required to the initial `3` to be misclassified as `7`, but also that they are not even required to mimic the targeted class (the manipulated `3`s evade detection without resembling a `7` at all). This confirms the vulnerability of standard learning algorithms to well-crafted attacks, and the need for more *secure* learning algorithms, as suggested by Barreno et al. (2006); Huang et al. (2011); Biggio et al. (2014).

4. Conclusions and future work

`AdversariaLib` is the first open-source library for the security evaluation of machine learning against carefully targeted-attacks. Currently, it implements the gradient-based evasion attack described in Biggio et al. (2013). As the library is modular and easy to extend, novel and more sophisticated attacks can be easily implemented, like *poisoning* (Biggio et al., 2012; Kloft and Laskov, 2010), as well as countermeasures and secure classifiers, hopefully, with the help of an emerging community of users and developers.

References

Marco Barreno, Blaine Nelson, Russell Sears, Anthony D. Joseph, and J. D. Tygar. Can machine learning be secure? In *Proc. ACM Symp. Information, Computer and Comm. Sec.*, ASIACCS '06, pages 16–25, New York, NY, USA, 2006. ACM.

B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrndić, P. Laskov, G. Giacinto, and F. Roli. Evasion attacks against machine learning at test time. In Hendrik Blockeel, Kristian Kersting, Siegfried Nijssen, and Filip Železný, editors, *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD), Part III*, volume 8190 of *Lecture Notes in Computer Science*, pages 387–402. Springer Berlin Heidelberg, 2013.

Battista Biggio, Blaine Nelson, and Pavel Laskov. Poisoning attacks against support vector machines. In John Langford and Joelle Pineau, editors, *29th Int'l Conf. on Machine Learning*, pages 1807–1814. Omnipress, 2012.

Battista Biggio, Giorgio Fumera, and Fabio Roli. Security evaluation of pattern classifiers under attack. *IEEE Transactions on Knowledge and Data Engineering*, 26(4):984–996, April 2014. ISSN 1041-4347.

L. Huang, A. D. Joseph, B. Nelson, B. Rubinstein, and J. D. Tygar. Adversarial machine learning. In *4th ACM Workshop on Artificial Intelligence and Security (AISec 2011)*, pages 43–57, Chicago, IL, USA, 2011.

M. Kloft and P. Laskov. Online anomaly detection under adversarial impact. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 405–412, 2010.

Y. LeCun, L. Jackel, L. Bottou, A. Brunot, C. Cortes, J. Denker, H. Drucker, I. Guyon, U. Müller, E. Säckinger, P. Simard, and V. Vapnik. Comparison of learning algorithms for handwritten digit recognition. In *Int'l Conf. on Artificial Neural Networks*, pages 53–60, 1995.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.